Cassini Solid State Recorder

A 1 ligh Capacity, Radiation-Tol erant 1 ligh-Performance Unit

K F Strauss, Jet Propulsion Laboratory*
G J Stockton, TRW, Inc., Electronics Systems & Technology Division*

Mission to Saturn

The Cassini spacecraft will soon begin a perilous journey into the depths of outer sp ace, exploring in-depth for the first, and perhaps the last, time the ringed-planet Saturn and its moons. One moon in particular, *Titan,* will be subject to particular scrutiny as it will receive an automated probe designed to sample the moon's at mosphere.
The Cassini mission consists of a 7 year cruise from launch to orbit insertion, 2-. years of which are within the region between Earth and Sun, using gravity assists from Earth and the planet Venus for injection into deep space.

Especiall y within the region of 1 Earth's orbit, the spacecraft and the Data it contains are subject to peril from Solar 1 Flares. in the orbit of Saturn, the craft will be bombarded with highly active, free protons trapped in an invisible. torus around the planet.

The Cassini Solid State Recorders - NASA's Preamble

Onboard the Cassini spacecraft, and the subject of this paper, arc two Solid State Recorders designed and built by TRW, inc., of Redondo Beach, California, for NASA's Jet I 'repulsion 1 aboratory. These recorders, which were the. first ones designed and selected by NASA for a space program, represent the new frontier and a new beginning for Spat.cbol'nc data storage.

This paper will compare the Cassini Solid State Recorders to moving tape recorders used for other J]'], missions, and then describe the advantages in going to solid slate.

1 Each SSR in the Cassini spacecraft has a capacity of 2 Gbits (data) at launch and are designed to allow for graceful degradation of data capacity as components age. and are irradiated to a guaranteed] End Of Mission capacity of 1.8 Gbits (data).

The SSRs were designed and hilt under a firm, fixccl-price. contract, which was awarded in 1992 through a competitive bid-and-review process[1].

Design Goals

The goal of the Solid State Recorder was to provide a sdid-slate replacement (not necessarily a drop-in replacement) for existing mechanical (moving) tape recorders. The capacity, size, and mass were to be comparable to its mechanical counterparts. '1'able 1,

* 4800 Oak Grove Drive, Pasadena, California 9 1 1 09
* One Space Park, 1OXIOI)(10 Beach, California 90278
[1] Jet Propulsion I aboratory contract no. 959291 , dated June 1, 1992, as amended

later, provides basic information on the Cassini SSR versus mechanical recorders used on Galileo, Voyager, Magellan, and Hubble.

A driving factor in the design and selection was the Bit Error Rate (BER), expressed to represent how many bits can be uncorrectable, and therefore corrupt, per every billion bits of data stored. The BER requirement for the SSR is:

Source	Period	Maximum Permissible Number of Bits in Error[2,3]
GCR (90% Worst Case), DCF @ 3 AU	7 days	3
	30 days	5
Saturnian-Trapped Protons	11 Orbits	0.8

Recorder Basics

The Solid State Recorder designed for Cassini is capable of storing 2 Gigabits of data. The storage medium is composed of 640 4-Mbit DRAMs manufactured by OKI Semiconductor of Japan and packaged and upscreened by TRW Components International Division of San Diego. The 640 DRAMs give the SSR a total data capacity of 2.56 Gigabits at launch, which include checksum bits reserved for the built-in Error Detection and Correction (EDAC).

The recorder features and highlights the ability to perform simultaneous data storage and access from multiple data ports and multiple areas of memory with no loss in processing capability or bandwidth. Data can be written into and read from any portion of memory, even the same portion, simultaneously. The combined read/write data rate is 4 Mbps with 2 Mbps per second maximum for any data port.

The most exploited feature of this recorder, and - which was new and unique at the time of contract - is the ability to segment the memory into 16 randomly sized "partitions." The partitions are treated as separate areas of memory contained within the whole. Each partition has its own set of independent "pointers" which are used as locators for where data is written into and read from. The uniqueness here is that each partition may be configured differently from any other - in terms of length, the number words per frame of data, and the "Mode" in which it operates. As a partition fills, certain safeguards are in place to guarantee that a partition's data does not "overflow" into an adjacent part of memory.

High-performance and optimum maintenance of the partitions is possible by two virtues of the SSR. One is that the partitions are built autonomously by the SSR. The interfacing Command and Data System (see Figure 1) merely sends a command requesting the size, the initial "mode," and the number of "frames" desired. If the sufficient memory exists,

[2] Per Gigabit for the Period specified.

[3] GCR = Galactic Cosmic Rays; DCF = Cassini Design Case Flare which approxiamtes Adams 90% - see JPL 94 flare data.

the command is immediately executed and ready (o receive important data. The second virtue is more complex, yet offers wide possibilities: Once created, and partition may be "moved" to anywhere else in memory by execution of certain commands. This allows, for the first time, the ability to predispose where certain, perhaps highly-critical, data is stored physically within the machine.

The advantage of being able to create multiple partitions in Memory becomes evident when each of up to 16 partitions can be set in any of three different modes of operation, separate from any other partition:

1. Read-Write to I 'nd,
2. Circular Buffer,
3. FIFO

Additionally, each partition can be:

a) Write Protected, such that extraneous or rogue attempts to write over data are blocked, and

b) 1 'ach memory cell comprising the partition can be checked for data integrity with a robust pattern without disturbing nearby parti tions.

Furthermore, each partition is sut>-divisible into "frames." A frame is a defined group of words within a partition. A partition may contain one frame or many frames. The idea of establishing data frames and the storage and retrieval of data in groups of frames opens wide the possibilities of expert data management and data safety.

Finally, the Cassini recorder features independent Command and Status channels which can also be accessed at a 2 Mb rate without impacting data bandwidth. The Status channel features two types of status switchable on command. One is a real-time reporting of the "active partitions". The other supplies over 400 words of infor mation about the status of the entire box.

As radiation, and life take their toll, certain areas of memory are expected to either dic outright, or become permanently corrupted so as to be of dubious quality. The goal of the Cassini SSR is to allow "graceful degradation" from 2 Gbits (data) at Beginning of Mission to 1.8 Gbits at End of Mission. Of 640 memory devices used in the SSR, 2 or 3 are expected to fail completely. Calculations performed to estimate. tile number of "stuck bits" (i.e. bits that have been impinged upon by ions of sufficient charge as to permanently damage them to assume forever and always one of the two possible logic states) at 61 K over the entire box. The combined effect is a potential loss of far less than 200 Mbits, thus affording ample margin,

Inner Basics of Recorder

Bulk Data Storage Medium

The Cassini SSR is based on what is now old technology: the 4 Mbit DRAM. The electrical design proposed by TRW caught a number of people by surprise in that it was based on DRAMs (Dynamic Random Access Memory) and not some other technology.

DRAMs, at that time, had always been the unwanted child of space-based storage technology as there was a preconceived notion of non-suitability.

Before the proposed design could be accepted, JPL independently embarked on a months-long course to validate the suitability of the OKI DRAMs. As JPL scrutinized the OKI parts for suitability, other parts - possibly suitable for deep space - were starting to emerge on the scene. JPL and TRW looked into parts from other vendors such as Micron, Toshiba, and Mitsubishi, to name just a few. The fared 1 BM 1 una-C's were just coming out of the foundry, and were reviewed, but rejected as suitable alternatives as there was no assurance at that time that IBM would produce them.

The OKI parts were subjected to Single Event Upset and Latchup testing at Brookhaven National Laboratory's Twin Tandem van de. Graaff generator and proton-induced Upset testing at the Harvard Cyclotron Laboratory to satisfy a contract requirement that single event upset characteristics be established for the flight-lot of parts.

The OKI parts used in the SSR (OKI part number MSM514400) were found to be latch-up immune (a Cassini requirement) to beyond 100 krads (Si). The have an LET in the range of 2 and are total dose tolerant to approximately 30 kRads before parametric shift occurs. Inside the recorder box itself, analysis has strewn that the piecepart total dose will be in the range of 3 krads and 1 krad inside the power converter hybrid[4].

LSI Logic rad hard gate arrays

The Cassini SSR contains 18 gate arrays of two designs: there are two Data Formatter gate arrays that sport upwards of 51,000 gates; anti sixteen Memory Controller gate arrays that have a gate, count in the 38,000 region, The gate arrays were designed using Synopsis design tools and a Zycad accelerator. The gate. arrays were manufactured by LSI Logic on their rad-hard LRH10K family to Class S requirements. The gate arrays utilize Boundary scan technology and feature a fault grading of over 99%, each.

Power Converter

Each SSR contains a pair of redundant power converters. The converters Were manufactured by Frequency Electronics, Incorporated, Long Island, New York. Each converter-half features a full-custom hybrid circuit. By placing all the critical control electronics inside a hybrid, a small, lightweight, highly-reliable power converter is achieved.

Either power converter (PCU) is capable of supporting all aspects of SSR operation independently. Each PCU contains a large capacitive hold-up bank that supports operation of the converter through periods of bus droop or fault.

Why the Move to Solid State?

As the speed and power of science processors increase, so does the desire for Scientists and Engineers to do more and more with (bc.in. And the business of space science is no different than any other. As new capabilities are discovered, or designed in, scientists and

[4] TRW memorandum 11620 JGWK.93-002A, 18 Feb 1993

engineers want more and more. In some instances, you have a single large spacecraft, such as Cassini, as the platform for a wide range of investigative instruments, or you have many spacecraft - each performing some small portion of the whole.. Currently, the pendulum is away from large multi-f faceted spat.ccraft and is in the realm of the small and the many. Cassini is, for now, the last of the large, all-encompassing, spacecraft.

Each science instrument has built-into - and an artifact of the basic design of it - a natural, optimal data rate at which it collects information, performs some type of translation or interpretation, and outputs it for use elsewhere in the universe.

Some instruments collect data very slow and methodically. Some collect data rapidly

With (ape recorders, the bandwidth of recording the data onto the storage medium (polymer tape) is directly related to the speed of the tape as it is transported across the record head.

In tape recorders, there is an inherent need to servo-loop the tape speed capstan drive to the incoming data clock. The slow instrument's data would be recorded using a slow tape speed whereas video camera information would require a very ilig,il-tape speed to achieve the necessary bandwidth. Galileo's lone tape recorder, for instance, goes up to 6 and one-half feet per second to accommodate a data rate of 800 kilobits per second. At that speed, lilt entire tape - over one-third *mile in length* - is used up in a mere 4 minutes. And hat urally, any increase in tape speed greatly exacerbates head wear; the tape Sill'faCC aCting like a high-speed rasp. Indeed, the use of tape and the amount of tape across a head is very tightly controlled; guarded much as the Crown Jewels.

Another method of increasing data rate in a moving tape recorder is to increase the number of heads recording on the tape at once, and that quite often requires a corresponding increase in the width of the tape.

Tape alignment across the record and playback heads is absolutely critical. As the number of heads increase the data-width for the tape approaches infinitesimal. A misalignment of tape would, at the very least, decrease the quality of data storage and retrieval. At the most, data could be lost.

The tape recorders for Galileo and Magellan use precision, optically controlled servo mechanisms to mechanically nudge the tape high or low as it soars pass the data heads.

The use of increasingly smaller tape. heads and gaps brings on another set of problems: Small gaps invite contaminant build-up. Many heads invite magnetic flux spill-over to other areas - tape capacity is lost bc.cause. of the need to include "guard band"

Whereas, in tape recorders the inertia of the tape drive mechanism has taken you wcii beyond the termination zone thus instilling an undesired inter-(iata gap of either old or missing information, in Solid State it is very possible to move a pointer to an exact word location.

Table I provides a comparative example for a few of the more famous spacecraft, Magellan, Galileo, and Hubble, versus the Cassini Solid State Recorder equipped spacecraft. Note the Tape Position Telemetry row: this indicates the granularity of ordering that the medium (tape/microcircuit) place a specific area strip of data at the data interface. None is bad, infinite is better.

Table 1: Spacecraft Data Recorder Comparative

SPACECRAFT / RAMETERS	GALILEO	MAGELLAN	HUBBLE	CASSINI SSR
PUT D ATA FORMAT	NRZ-L	NRZ-L	NRZ-L	NRZ-L
PUT CLOCK SOURCE	INTERNAL, 2x	EXTERNAL 2X	EXTERNAL, 2X	EXTERNAL 1X
PUT D ATA RATE, /SEC	7.68, 28.8. 115.2, 403,2, 304.6	7.7, 806	4- 1024 (4 RATES)	DC to 2 Mbps
CORD TIME	8 HR slowest 4 MIN fastest	65 HRS - 37 MIN	10.4 HRS - 19.5 MIN	17 MIN at Max data rate
CORD SPEED, IPS	0.738, 2.77, 11.07, 38.76, 77. s4	0.4, 42.2	1.3 - 41	NA
TAL STORAGE BITS	9×10^8	1.8×10^9	1.5×10^8 (4 Kbps) 1.2×10^9 (32 . 1024 Kbps)	2×10^9 data (BOM) 1.8×10^9 (gtd EOL)
TA TRACKS	4	4	2	16 partitions
TPUT DATA FORMAT	NRZ-L	NRZ-L	NRZ-L	NRZ-L
TPUT CLOCK URCE	EXTERNAL	EXTERNAL	EXTERNAL (2x)	EXTERNAL 1 X
TPUT DATA RATE, /SEC	7.68, 19.2, 57.6, 100.8	15, 26?	1024	DC to 2 Mbps
PRODUCE TIME	8 HR slowest 37 MIN fastest	4.3 HRS - 1.87 HRS	?0 MIN	17 MIN at Max data rate
PRODUCE SPEED, 11'S	0.738, 1.846, 5.54, 9.69	6.0, 14	41	NA
ERR OR RATE	5 in 10^6	1 in 10^6	1 IN 10^8	Solar Flare: 3 bits/1 0"/ week 5 bits/10^9/month Saturn trapped pr otons: 0.8 bits/Gbit/hour
PUT JITTER,	±0.1%	±0.1%	±0.1%	None
ERATING VOLTAGE,	+30V ±5%	+28 ± 1	+28 ±4	22 to 35 37 111s bus-fault hold-up
CORD POWER, \v	9.0 - 17	22	28	9
PRODUCE P0\\'ItH, W	13 - 14.2	22	28	9
PE LENGTH, (FEET)	1850	1968	2068	NA
MMAND INTERFACE	DISCRETE CMOS	DISCRETE CMOS	16 BIT SERIAL (TTL)	16 bit Current Mode (HS245/HS246)
PE POSITION EMETRY	NONE	NONE	127 INCREMENTS	Infinite pointer control
IENSIONS, INCHES	7 X 9 X 12.8	7 X 9 X 15	7 X9X12.8	7X 7.185 x 10.44
UME, CU. INCHES	806	)s0	306	826
IGHT, POUNDS	19.8	22	22	30

'1'able. Ishows that the Cassini Solid State Recorder has wider bandwidth and features faster and more precise across to specific data with less power and less weight than its mechanical predecessors.

'1 'wo major advantages readil y apparent as we look at the move to solid state.:

1) While data rate affects tape speed and thus wear out of the (ape and heads, no such mechanism exists in Solid Slate. And,
2) Only solid state designs permit for different record and playback d ata rates simultaneousl y

Need to optimize data storage for device

To optimize the use of downlink resources, current-design spacecr aft literall y bundle, or "packetize," the gathering of information from many resources and send down what is effectivel y a cont inuous stream of informat ion. The information contained in each individual packet is identified by an ingrained header. When information from one instrument is not avail able, instead of dead air time, a 'packet' of information is grabbed from somewhere else, and placed on t he ai rwaves.

Just as each instrument is optimized for certain data rates, the amount or length of data is also optimized. Where one instrument provides information in as little of, say, 12 words of data, mother may requite hundreds. In Solid State recorders, it is possible. - with no change in mechanics - to optimize each "track, " or *partition*, of memory to the specific instrument being stored at the time. This gives a unique capability not found in tape recorders: it is instantly possible to locate individual fields, frames, or even bits of data for replay or analysis.

History of Recorders

Moving tape

Tape Recorders are still in production today. Recorders are being proposed in the range of Terabit capacity. The drawback 10 this is the large weight (hundreds of pounds) and power (hundreds of watts) penalty that comes with such a large machine.

Tape recorders present a control problem in spacecraft - the moments of inertia change as a [ape stalls and stops, speeds up 01' slows down. Only certain designs of 'coaxial' (ape recorders cancel the effects of tape mass transferring from one reel to another. Springs are used to balance tension and motion, with every flexure being one step close.r to failure.

occasionally, there is the need to locate certain blocks of data. 1 n a moving tape. recorder often limes that data falls between the small slits in the position encoder anti the tape carefully jogged past the heads. But sometimes that is not good enough - positioning to the bi[-level has been attempted seve1 al times in recent history.

Solid State

Solid state memory has existed for a long time: core. While it is inherently rad-hard and upset proof[5], the capacity and speed did not exist for today's spacecraft. Core is expensive to manufacture - threaded by hand using wires much finer that a human hair - and comparatively very slow.

Just over the past few years has the per-chip capacity of solid state memory increased to a point as to be feasible for mass data storage in space. And the Cassini solid state recorder, we are proud to say, was the very first.

Design of the Cassini Units

When the decision was made to go solid state for Cassini, we were, at the time, examining what specialized circuitry would be required to accommodate either the spare Galileo or Magellan tape recorders. The spare Magellan recorder, with its higher 1.8 Gbit capacity became the baseline for the mission and thus proved the driving, force in deciding available data capacity for Cassini.

Any solid state recorder would have to be competitive in power, size, and mass with its mechanical forefathers. With power developed by three thermonuclear generators, power aboard Cassini is at an absolute premium.

Where as a tape-based recorder can withstand power loss, solid state devices, in general[6], have no such luxury. 1 lere, the loss of power can cause the complete loss of science data and navigational and control programs. Thus the Cassini design was specified to be tolerant Of power bus faults Of 37 milliseconds, whether the bus voltage abruptly terminate.d (blackout) or slowly decayed (brownout). The ('assini SSR power converter contains within a bank of hold-up capacitors to hold power and interfaces viable for bus-fault conditions.

The immediate downside to the selection of DRAMs for data storage in deep space is that they are highly subject to upsets causal by solar flare, cosmic rays, heavy ions, or, in Cassini's case, the dangerous trapped-proton belt of Saturn. "1'0 guard against data upset and still be bounded by reasonable error correcting scrub rates (which drives power requirements) the Cassini SSR is heavily shielded on all sides by an equivalent of 500 roils of aluminum. Wc have taken a "nit" in mass 10 ensure data integrity.

The Cassini SSR interfaces with two Command Data Subsystems - a principal of block redundancy. (see Figure 1, below.) livery aspect of SSR operation is controllable or accessible via either port with absolutel y no need (o reload cent rolling parameters as cont rol is handed-off between Ports A and B. Further, the design of the SS1< permits operation, in a prescribed manner, immediately following power-up.

[5] NDRO (Non-Destructive ReadOut) core
[6] excluding, 01 course, core, flash, and EEPROM

Representative CDS-SSR cross-string interconnect

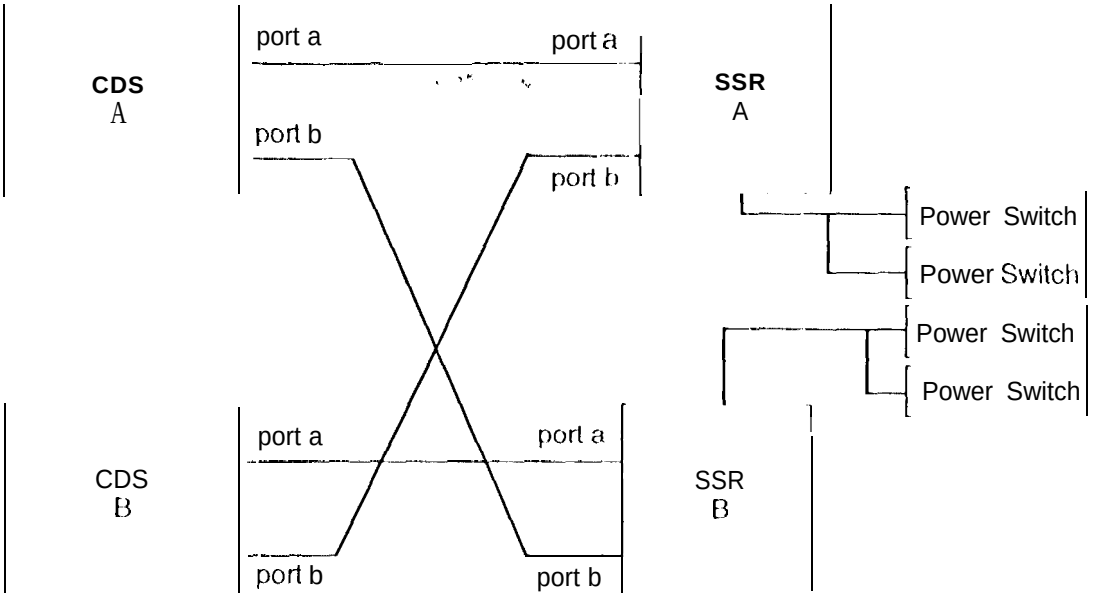

Figure 1: CDS- SSR Interconnect Diagram

Diagnostics

With each SSR containing 640 individual DRAMs and with each DRAM having a capacity of over 4 million bits of data, and with 4 device transistors assigned to each bit, one can see by the pure math alone, that there is a possibility that something is bound to go wrong somewhere[7]. To that end, the SSR contains a 6-pattern self-check that can be specified over ally particular area of memory, very small or very large. The results Of the self- check are stored and reported for ground analysis. In order to absolutely minimize the possibility of lost data, the SSR dots nothing that would affect data autonomously - it only alters memory configurations on CDS or ground commands.

Data Integrity

The SSR contains built-in EDAC executing the (39,32) Hamming SECDED (Single Error Correction, Double Error Detection) code. The entire memory is scrubbed for errors every 538 seconds. Additionally, data read from the SSR is scrubbed to ensure correctness. Single bit errors are corrected, double bit errors are detected. The Single bit/double bit error counts are reported for every Megaword of memory. Ground analysis is then used to determine if a failure trend exists. Ground controllers can set, via command, that certain areas of memory not be used for' data storage. The "off" areas are taken into account for when a new partition is built.

[7] The SSRs have a calculated reliability of 98.6%.

Implementation of Requirements

TRW implemented data and control entirely by hard-silicon: Custom Radiation-hardened Gate Arrays. In this case, each SSR contains a total of 18 ASICs, comprised of two Data Formatters and 16 Memory Controllers. One Data Formatter is primarily connected to each Port and both Data Formatters have access to all sixteen MCs. (See Figure 2) The ASICs are all based on the 1S1 Logic LRH1 OK family.

CASSINI SOLID STATE RECORDER BLOCK DIAGRAM

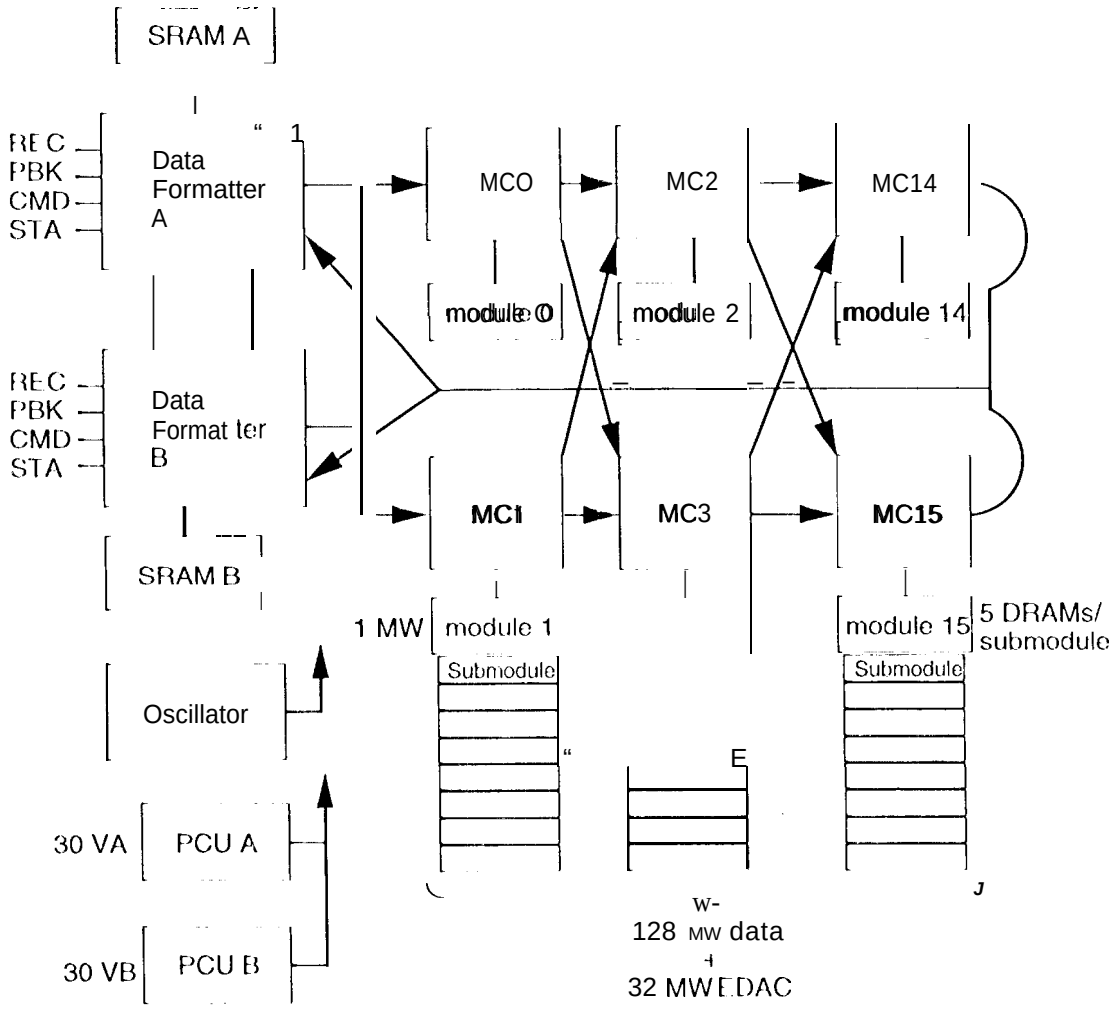

Figure 2: Internal SSR Interconnect

The Data Formatters (DFs) microcontrol each operation of the SSR. The DFs receive data and commands from the CDS and output data and status to it. The DFs act as the sole interface between data and the outside world.

The])] 's decode commands and passes them downstream to the Memory Controllers. The Memory Controllers (MCs) act as the interface between the data and the active Data Formatter. The DF will translate *logically* based information from the CDS (e.g. Word 23 of partition 6) to a physical address, such as Word 07FEEDBA. The Data Formatter acts as the master timing element within the CDS, ordering when DRAM-required refreshes and error correcting scrubs of data is to take place; with the Memory Controllers actually performing the. work.

Each MC is responsible for 8 M Words of data with start and stop addresses of its block of memory set via printed wiring board jumpers. The MC watches the internal SSR bus for address header information. If it is within the realm of its address block, the MC will grab the header and act accordingly. The MC provides local DRAM refresh timing and local EDAC error scrubbing.

Control and initial placement of the. various data pointers is performed (calculated) autonomously by the Data Formatter and physical addresses are stored in rad-hard IBM 32K x 8 SRAM. It is possible, via diagnostics routines to entirely predetermine the size and placement of each partition by pre-loading the SRAM with carefully calculated parameters.

As shown in Figure 2, the data chain internal to the SSR consists of two parallel-serial paths. Data and commands are passed from MC to MC. A unique aspect of the Cassini SSR design is the "I'M-patented, 3-pint architecture. The failure of one Memory Controller in the data path does not prevent the retrieval or storage of data throughout the rest of the recorder. The SSR has built-in the ability to bypass single point failures. This architecture provides two levels of fault tolerance with a third provided by redundant Data Formatters and redundant I/O polls.

Each MC has two input ports and one output port, hence the term 3-port architecture. The two input ports allow the MC (or *module*) to receive data and commands from either of two directions, The output port fans out in two directions, which allows a failed module to be completely bypassed. Even multiple-module failures are tolerable. A built-in "ping" test allows the Data Formatter to ascertain the health of the entire data chain and for an MC to determine the health of its neighbors.

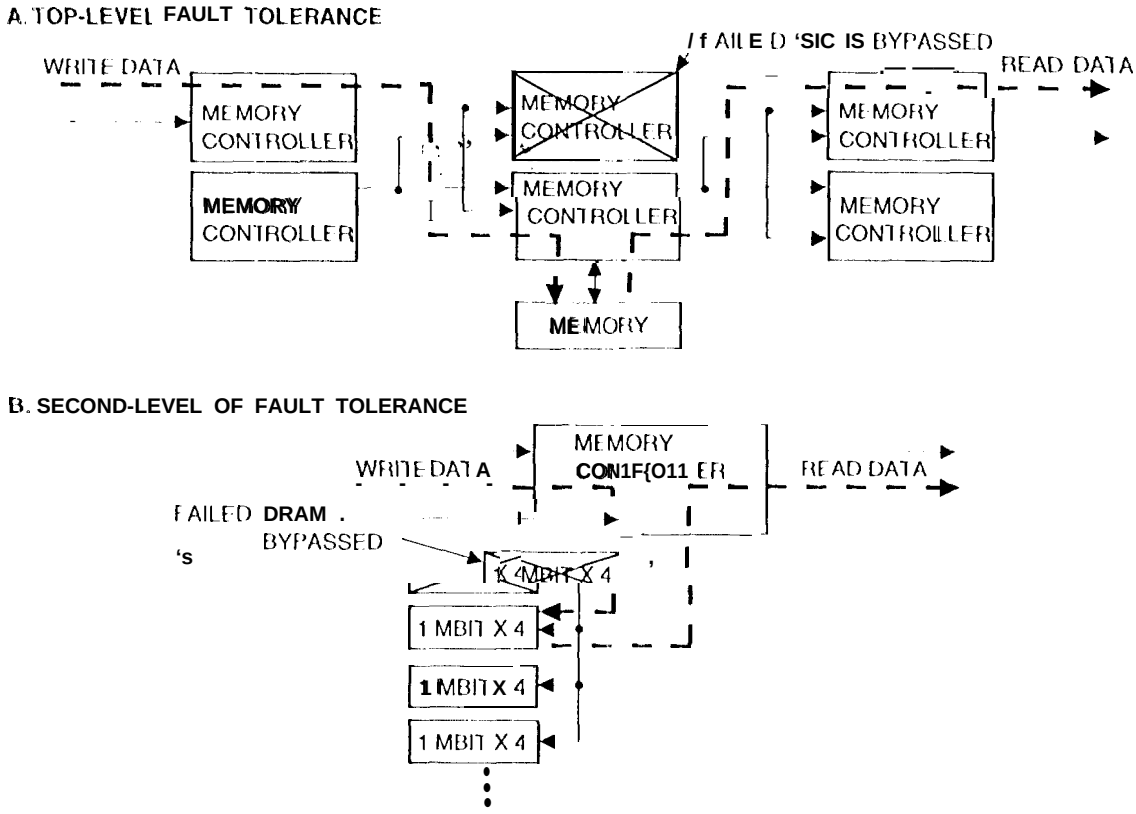

Figure 1.1-2. Method of Bypassing Failed AS I Cs/DRAMs for Fault Tolerance

This is a very robust design technique.

Partitions and Frames

The strong forte of the Cassini SS1< is the built-in capability to divide the bulk data memory into 16 randomly sired partitions. [Each partition, in turn, is divisible into "frames," where a frame represents an optimal block of data.] Each partition has its own "set" of Pointers used to indicate where data is lining written into (the Record pointer) and read from (the Playback pointer). Thus there are up to 16 pair of data pointers. Each pointer is independent from any other; however, only one Record pointer and only one Playback pointer may be active at any one time. The activation of a pointer is done by a simple command.

It is possible that any partition have both the active Record and Playback pointers; and it is entirely possible that one partition have the Active Record pointer and another partition have the Active Playback pointer. Thus it is possible to read and write separate area of memory at the same time.

Each partition is sub-divisible into "frames." A frame is a defined group of words within a partition. A partition may contain one frame., or many frames. The idea of establishing data frames and the storage and retrieval of data in groups of frames opens wide the possibilities of expert data management and data safety.

1)ata is written into and read out of the SSR serially. Control of the active data channel is done by 1 :nable/Ready handshaking. The SS1< expects data to be transacted in groups of frames for each partition. Should a control signal be deasserted on a non-frame boundary that is an indication that something has gone wrong, either not enough data was passed, or perhaps too much. In either case, the SSR sets a Status alert flag that the last transaction was somehow corrupted. In order to protect just written data, which may actually be good, the SSR will automatically move the Record Pointer ahead to the next frame boundary. In order to allow a quick re-read of data, the SSR moves the affected P layback Pointer backwards to the start of the current frame.

The SSR operates each partition independently from any other. Each partition may be operated in any one of three distinct modes of operation, each individually settable and interchangeable for all partitions of memory.

Mode l,l<c:i(l-Write-to-l{l~(J. This mode emulates a reel-to-reel tape recorder in that memory is accessed logically from beginning to end and when the end of the partition block is reached, data activity ceases.

Mode 2, Circular Buffer. In this mode, there is no "beginnin g and end" of a partition in memory. The last logical address is seamlessly looped to be electrically adjacent to the first, thus forming an apparent, virtual continuous circle of I1KXI101Y. This mode is particularly useful for storage. of information that is frequently overwritten and for which Icli:il]cci s]ll:i(ico llttlc''latcst data available."

Mode 3, FIFO. In this mode, as in Mode 2, there is no beginning nor end of memory addresses; it also forms a continuous circle. 1 however, this mode is unique in that the individual read and write pointers arc prohibited from passing each other in order to protect unwritten or unread data.

One additional mode, not intended to be used except by ground intervention, is referred to as J)irect Memory Address (DMA). in DMA any area of memory, down to an exact word is immediately accessible and alterable. The DMA mode permits intensive testing of a memory device to assist in trend analysis,

"Slice" based architecture

In that the SSR is "segmental" is very unique. Memory can be added or subtracted with the relatively simple installation or removal of "slices" of memory; a major attribute of using a serial-basc(i architecture with local memory control.

The Cassini SSR consists of 4 mechanical "slims." The Power Converter counts as one slice. The next slice inward - the I/O slice - contains printed wiring boards upon which are mounted the two data formatters and four Memory Controllers. The remaining, two slices are termed Memory Slice.s. Each contains six additional MCs and 48 MWords of data storage. Ail the slices are interconnected via a set of custom wiring harnesses.

CONCLUSION

This paper has discussed the Cassini Solid State Recorders and the reason why a solid state design was chosen over mechanical, moving tape. This paper has shown the virtues of a solid state design versus a moving tape machine.

Why the design of the Cassini SSR is now old, it offered for the first time, unique data control capabilities which are now implemented as standard in almost all new SSR designs today.

Bibliography

Cassini Solid State Recorder Functional Requirements, FM515834C, December 21, 1992, Jet Propulsion Laboratory.

Equipment Specification for Cassini Solid State Recorder, EQ4-04 36D, 1/27/95, TRW, inc., Electronics Systems & Technology Division

Strauss, K., SSR User's Guide, ver 2.0, D-12777, 6/21/95, Jet Propulsion Laboratory

Swift, G., Single Event Rates for the Cassini Solid State Recorder. Interoffice Memorandum GMS-507-E-XXX-95, May 17, 1995, Jet Propulsion Laboratory

CASSINI SOLID STATE RECORDER BLOCK DIAGRAM

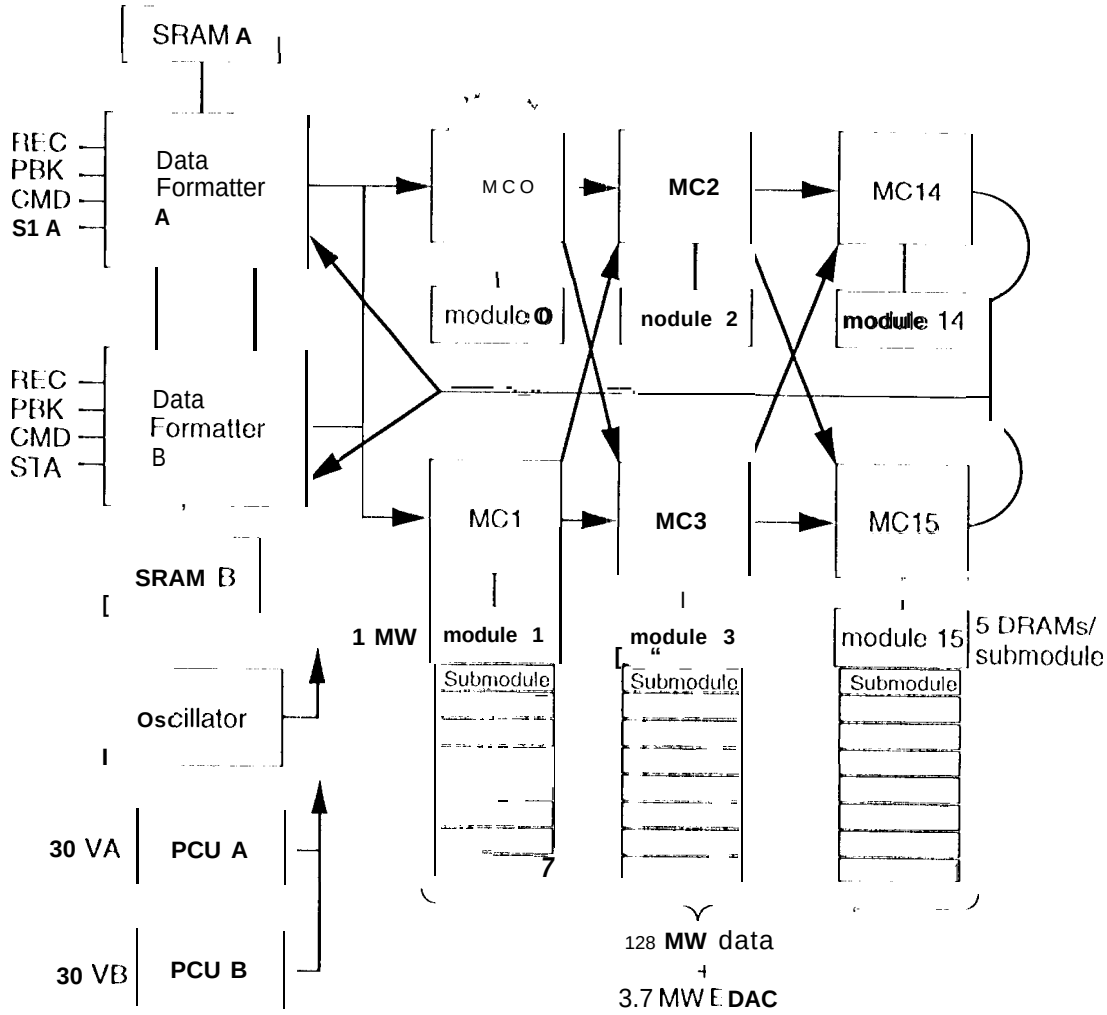

A TOP-LEVEL FAULT TOLERANCE

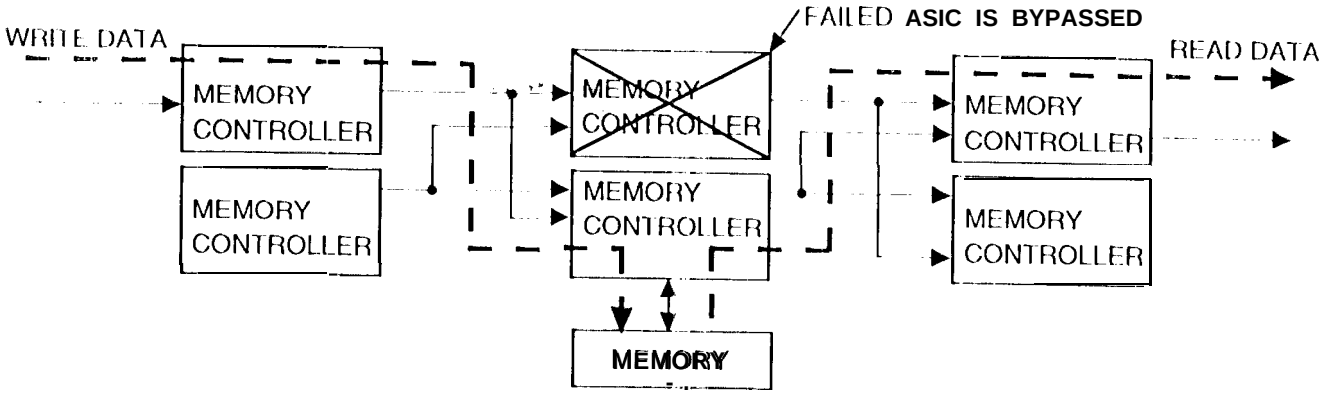

B. SECOND-LEVEL OF FAULT TOLERANCE

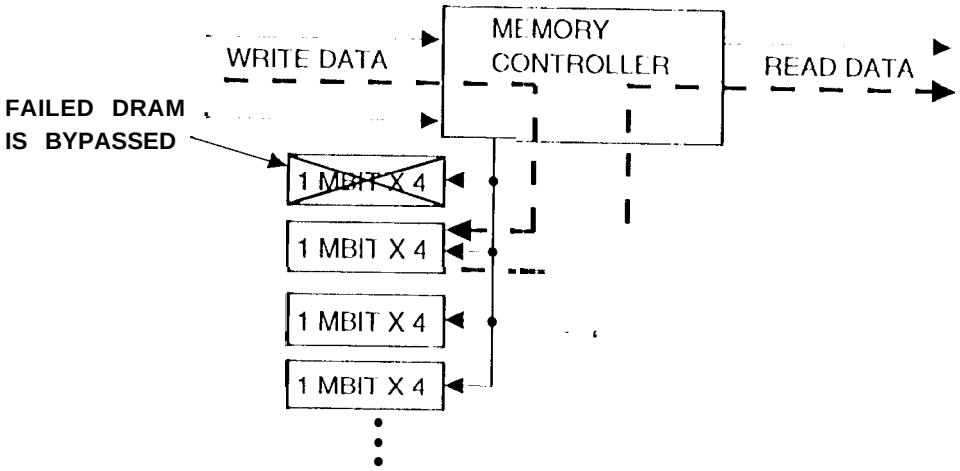

Representative CDS-SSR cross-string interconnect